

ALGEBRAIC AND MODEL THEORETIC METHODS IN CONSTRAINT SATISFACTION

MICHAEL PINSKER

ABSTRACT. This text is related to the tutorials I gave at the Banff International Research Station and within a “Doc-course” at Charles University Prague in the fall of 2014. It describes my current research and some of the most important open questions related to it.

1. OVERVIEW

A *function clone* is a set \mathcal{C} of finitary functions on a set D which is closed under composition and which contains all projections. More formally,

- whenever $f \in \mathcal{C}$ is n -ary, and $g_1, \dots, g_n \in \mathcal{C}$ are m -ary, then the m -ary function $f(g_1, \dots, g_n)$ defined by

$$(x_1, \dots, x_m) \mapsto f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m))$$

is an element of \mathcal{C} ;

- for all $1 \leq k \leq n < \omega$, \mathcal{C} contains the k -th n -ary projection $\pi_k^n: D^n \rightarrow D$, uniquely defined by the equation $\pi_k^n(x_1, \dots, x_n) = x_k$.

There are two main sources of function clones:

- the term operations of any algebra \mathfrak{A} form a clone, the *term clone* of \mathfrak{A} (and in fact, every clone is of this form);
- the set of all operations which *preserve* a given relational structure Γ form a clone, called the *polymorphism clone* of Γ (certain clones, the *topologically closed clones*, are of this form).

The first source of function clones makes them an object of primary interest in universal algebra, since many properties of an algebra, such as its subalgebras and congruences, only depend on its term operations. The second source links them with relational structures, and in particular, as we will see, with certain questions in complexity theory.

This topic of this text are function clones over a countably infinite set and their applications in complexity theory. While the investigation of *all* such clones is not very promising since in the general setting, hardly any positive structural results could be expected (cf. for example [26]), research on clones which are “sufficiently rich” has proven extremely fruitful in recent years [11, 12, 5, 36]. We are interested here in function clones which are rich in the sense that they contain a rather large permutation group: a permutation group on a countably infinite set D is called *oligomorphic* iff its componentwise action on D^n has only finitely many orbits, for all $n \geq 1$ (cf. [23]). Function clones containing an oligomorphic permutation group, referred to as *oligomorphic clones* [5], have been shown to enjoy many properties of function clones on finite sets. For example, they satisfy a topological variant of Birkhoff’s HSP theorem; moreover, they encode the complexity of certain computational problems, so-called *constraint satisfaction problems (CSPs)*, and indeed have proven to be a valuable tool

in the study of the complexity of such problems in what is called the *algebraic approach* to CSPs. Oligomorphic function clones encode a much larger class of CSPs than function clones over finite sets [4], and yet many tools from the finite carry over to the oligomorphic setting.

Of particular importance to us will be oligomorphic clones which arise from *homogeneous* relational structures in a finite language, and some of our methods rely on *Ramsey theory* and on connections of Ramsey-type theorems with topological dynamics. While the original motivation for studying function clones comes from universal algebra, and later and independently from constraint satisfaction problems, their study therefore also involves tools and concepts from model theory, combinatorics, and topological dynamics.

2. THE STATE OF THE ART

Birkhoff's theorem for oligomorphic clones. I will start by recalling the finite version of Birkhoff's HSP theorem. An *algebra* is a structure with a purely functional signature. The *clone of an algebra* \mathfrak{A} with signature τ , denoted by $\text{Clo}(\mathfrak{A})$, is the set of all functions with finite arity on the domain A of \mathfrak{A} which can be written as τ -terms over \mathfrak{A} . More precisely, every abstract τ -term t naturally induces a finitary function $t^{\mathfrak{A}}$ on A , and $\text{Clo}(\mathfrak{A})$ consists precisely of the functions of this form.

Let $\mathfrak{A}, \mathfrak{B}$ be algebras of the same signature τ . The assignment ξ from $\text{Clo}(\mathfrak{A})$ to $\text{Clo}(\mathfrak{B})$ which sends every element $t^{\mathfrak{A}}$ of $\text{Clo}(\mathfrak{A})$ to $t^{\mathfrak{B}}$ is a well-defined function if and only if for all τ -terms s, t we have that $s^{\mathfrak{B}} = t^{\mathfrak{B}}$ whenever $s^{\mathfrak{A}} = t^{\mathfrak{A}}$. In that case, ξ is in fact a surjective *clone homomorphism*, and we then call ξ the *natural homomorphism* from $\text{Clo}(\mathfrak{A})$ onto $\text{Clo}(\mathfrak{B})$. In general, a *clone homomorphism* is a function $\sigma: \mathcal{C} \rightarrow \mathcal{D}$, where \mathcal{C}, \mathcal{D} are clones (possibly acting on different base sets), which sends functions in \mathcal{C} to functions of the same arity in \mathcal{D} , every projection in \mathcal{C} to the corresponding projection in \mathcal{D} , and which preserves composition, i.e., $\sigma(f(g_1, \dots, g_n)) = \sigma(f)(\sigma(g_1), \dots, \sigma(g_n))$ for all n -ary $f \in \mathcal{C}$ and all m -ary $g_1, \dots, g_n \in \mathcal{D}$, for all $n, m \geq 1$ (cf. [16]).

When \mathcal{C} is a class of algebras with common signature τ , then $\text{P}(\mathcal{C})$ denotes the class of all products of algebras from \mathcal{C} , $\text{P}^{\text{fin}}(\mathcal{C})$ denotes the class of all *finite* products of algebras from \mathcal{C} , $\text{S}(\mathcal{C})$ denotes the class of all subalgebras of algebras from \mathcal{C} , and $\text{H}(\mathcal{C})$ denotes the class of all homomorphic images of algebras from \mathcal{C} (when defining these operators, we consider algebras up to isomorphism). A *pseudovariety* is a class \mathcal{V} of algebras of the same signature such that $\mathcal{V} = \text{H}(\mathcal{V}) = \text{S}(\mathcal{V}) = \text{P}^{\text{fin}}(\mathcal{V})$, i.e., a class closed under homomorphic images, subalgebras, and finite products; the pseudovariety *generated* by a class of algebras \mathcal{C} (or by a single algebra \mathfrak{A}) is the smallest pseudovariety that contains \mathcal{C} (contains \mathfrak{A} , respectively). For *finite* algebras, Birkhoff's HSP theorem takes the following form (see Exercise 11.5 in combination with the proof of Lemma 11.8 in [22]).

Theorem 1 (Birkhoff [3]). *Let $\mathfrak{A}, \mathfrak{B}$ be finite algebras with the same signature. Then the following three statements are equivalent.*

- (1) *The natural homomorphism from $\text{Clo}(\mathfrak{A})$ onto $\text{Clo}(\mathfrak{B})$ exists.*
- (2) $\mathfrak{B} \in \text{HSP}^{\text{fin}}(\mathfrak{A})$.
- (3) \mathfrak{B} *is contained in the pseudovariety generated by \mathfrak{A} .*

When \mathfrak{A} and \mathfrak{B} are of arbitrary cardinality, then the equivalence of (2) and (3) still holds; however, if one wants to maintain equivalence with item (1), then another version of Birkhoff's theorem states that one has to replace finite powers by arbitrary powers in the second item, that is, one has to replace $\text{HSP}^{\text{fin}}(\mathfrak{A})$ by $\text{HSP}(\mathfrak{A})$; the third item has to be adapted using the

notion of a *variety* of algebras, i.e., a class of algebras of common signature closed under the operators H, S and P.

It recently turned out that one can prove a similar theorem with finite powers for algebras \mathfrak{A} on a countably infinite domain whose clone $\text{Clo}(\mathfrak{A})$ is oligomorphic – we call such algebras *oligomorphic* as well. To this end, one has to see function clones not only as algebraic, but also as topological objects. On any set D , there is a largest function clone \mathcal{O}_D : the clone of all finitary operations on D . The “function space” \mathcal{O}_D carries a natural topology, namely the topology of pointwise convergence, with respect to which the composition of functions is continuous. A basis of open sets of this topology is given by the sets of the form

$$\{f: D^k \rightarrow D \mid f(a_1^1, \dots, a_k^1) = a_0^1, \dots, f(a_1^n, \dots, a_k^n) = a_0^n\}.$$

In fact, similarly to the Baire space $\mathbb{N}^{\mathbb{N}}$, \mathcal{O}_D then becomes a Polish space (cf. for example [16]). As a subset of \mathcal{O}_D , every function clone \mathcal{C} on D inherits this topology, and hence carries a topological structure in addition to its algebraic structure given by the equations which hold in \mathcal{C} . We denote the topological closure of a function clone \mathcal{C} in \mathcal{O}_D by $\overline{\mathcal{C}}$.

It is not hard to see that all algebras in the pseudovariety generated by an oligomorphic algebra are again oligomorphic. The following is the topological variant of Birkhoff’s theorem for oligomorphic algebras.

Theorem 2 (Bodirsky and Pinsker [11]). *Let $\mathfrak{A}, \mathfrak{B}$ be oligomorphic or finite algebras with the same signature. Then the following three statements are equivalent.*

- (1) *The natural homomorphism from $\overline{\text{Clo}(\mathfrak{A})}$ onto $\overline{\text{Clo}(\mathfrak{B})}$ exists and is continuous.*
- (2) *$\mathfrak{B} \in \text{HSP}^{\text{fin}}(\mathfrak{A})$.*
- (3) *\mathfrak{B} is contained in the pseudovariety generated by \mathfrak{A} .*

Note that Theorem 1 really is a special case of Theorem 2, since the topology of any function clone on a finite set is discrete, and hence the natural homomorphism from the clone of a finite algebra to that of another algebra is always continuous.

Applications to constraint satisfaction problems. Let us now turn to applications of oligomorphic function clones to computational complexity problems. Every relational structure Γ in a finite language defines a computational problem, called the *constraint satisfaction problem* of Γ and denoted by $\text{CSP}(\Gamma)$, as follows: input of the problem is a primitive positive sentence ϕ in the language for Γ , i.e., a sentence of the form $\exists x_1, \dots, x_n (\phi_1 \wedge \dots \wedge \phi_m)$ where ϕ_1, \dots, ϕ_m are atomic formulas; the problem is to decide whether or not ϕ holds in Γ . An instance of this problem therefore asks about the existence of elements of Γ satisfying a given conjunction of atomic conditions. The structure Γ is called the *template* of the problem, and can be finite or infinite. We will later see how infinite templates can model natural computational problems, and refer also to [4] for an abundance of examples. We remark that $\text{CSP}(\Gamma)$ is often presented in the form of a *homomorphism problem*, which is easily seen to be equivalent: in this formulation, the input is a *finite* structure Ω in the language of Γ (which can still be finite or infinite), and the question is whether or not there exists a homomorphism from Ω into Γ .

To every relational structure Γ , one can assign a function clone on the domain of Γ as follows. A *polymorphism* of a structure Γ is a homomorphism from Γ^k to Γ for some finite $k \geq 1$; the *polymorphism clone* $\text{Pol}(\Gamma)$ of Γ is the set of all polymorphisms of Γ . It is easy to see that $\text{Pol}(\Gamma)$ is a function clone which is closed in the pointwise convergence topology

described above, and in fact, the closed function clones are precisely the function clones of the form $\text{Pol}(\Gamma)$ for a relational structure Γ .

For finite relational structures Γ , the complexity of $\text{CSP}(\Gamma)$ depends, up to polynomial time, only on $\text{Pol}(\Gamma)$ (cf. [21, 20]); this fact is the basis of the approach to constraint satisfaction via clones. The same is true for structures with oligomorphic polymorphism clones [10]. But which structures have oligomorphic polymorphism clones? The answer can be found in a classical theorem of model theory, the theorem of Engeler, Svenonius, and Ryll-Nardzewski (see e.g. the textbook [28]). A countable structure Γ is called ω -categorical iff all countable models of the first-order theory of Γ are isomorphic to Γ . Now the theorem states that the automorphism group $\text{Aut}(\Gamma)$ of a countable structure Γ is oligomorphic if and only if Γ is ω -categorical. It follows that the polymorphism clone of a countable structure Γ is oligomorphic if and only if Γ is ω -categorical.

Theorem 3 (Bodirsky and Nešetřil [10]). *Let Γ, Γ' be ω -categorical structures in finite relational languages which have the same domain. If $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma')$, then $\text{CSP}(\Gamma')$ is polynomial-time reducible to $\text{CSP}(\Gamma)$.*

As a consequence, for ω -categorical structures Γ the complexity of their CSP is still up to polynomial time encoded in their polymorphism clone, i.e., if $\text{Pol}(\Gamma') = \text{Pol}(\Gamma)$, then $\text{CSP}(\Gamma)$ and $\text{CSP}(\Gamma')$ are polynomial-time equivalent.

The theory of the algebraic approach to CSPs goes much further, which brings us back to Birkhoff's HSP theorem. For a structure Γ , we call any algebra on the domain of Γ whose functions are precisely the elements of $\text{Pol}(\Gamma)$ indexed in some arbitrary way a *polymorphism algebra* of Γ . It can be shown that if Γ and Γ' are finite structures in a finite relational language, and a polymorphism algebra \mathfrak{B} of Γ' is contained in the pseudovariety of a polymorphism algebra \mathfrak{A} of Γ , then $\text{CSP}(\Gamma')$ is polynomial-time reducible to $\text{CSP}(\Gamma)$ [21, 20]. By Birkhoff's theorem, this is the case iff the natural homomorphism from $\text{Clo}(\mathfrak{A})$ onto $\text{Clo}(\mathfrak{B})$ exists. One then sees that this is the case if and only if there exists a surjective clone homomorphism from $\text{Pol}(\Gamma)$ onto $\text{Pol}(\Gamma')$. Hence, the complexity of the CSP of a finite relational structure Γ only depends on the abstract structure of the clone $\text{Pol}(\Gamma)$. Similarly to abstract groups, *abstract clones* can be formalized as multi-sorted algebras equipped with composition operations as well as with constant symbols for the projections (cf. , for example, [26] or [16]), but one can avoid this technicality: in practice, it is enough to know that abstract clones simply encode the equations which hold between its functions, or more precisely, it is enough to know that clone homomorphisms as defined above are precisely the structure preserving maps between those objects.

Using the topological generalization of Birkhoff's theorem, one can show the following for ω -categorical structures (this is a simplified version; for a stronger formulation see [11]).

Theorem 4 (Bodirsky and Pinsker [11]). *Let Γ, Γ' be finite or ω -categorical structures in a finite relational language. If there exists a surjective continuous clone homomorphism from $\text{Pol}(\Gamma)$ onto $\text{Pol}(\Gamma')$, then $\text{CSP}(\Gamma')$ is polynomial-time reducible to $\text{CSP}(\Gamma)$.*

In particular, for ω -categorical structures the complexity of their CSP is still up to polynomial time encoded in their polymorphism clone, seen as an abstract clone together with the topology on the functions. In analogy to topological groups, we call such objects *topological clones* [16]. More precisely, if two ω -categorical structures Γ, Γ' have polymorphism clones

which are isomorphic as topological clones (i.e., via a bijection which is a clone homomorphism, whose inverse is a clone homomorphism, and which is a homeomorphism), then their CSPs are polynomial-time equivalent.

A class of ω -categorical structures for which the CSP is of particular interest are structures with a first-order definition in a homogeneous structure in a finite language. A structure Δ is called *homogeneous* iff any isomorphism between finitely generated substructures of Δ extends to an automorphism of Δ (some authors call this notion *ultrahomogeneity* to distinguish it from related concepts of homogeneity). A *reduct* of a structure Δ is a relational structure on the same domain each of whose relations can be defined in Δ by a first-order formula without parameters. Countable homogeneous structures in a finite relational language are ω -categorical, and reducts of ω -categorical structures are ω -categorical as well, and hence fall into our context (cf. the textbook [28]).

When Γ is the reduct of a homogeneous structure in a finite language, then $\text{CSP}(\Gamma)$ models a certain type of problem about finitely generated structures, as we will outline in the following. Homogeneous structures can be seen as generic objects, called *Fraïssé limits*, representing so-called *Fraïssé classes* of finitely generated structures. A *Fraïssé class* is a class \mathcal{C} of finitely generated structures in a fixed countable language closed under isomorphism and induced substructures which satisfies the joint embedding property, i.e., for all $\Omega_0, \Omega_1 \in \mathcal{C}$ there is $\Omega_2 \in \mathcal{C}$ such that Ω_0, Ω_1 embed into Ω_2 , and the *amalgamation property*, i.e., for any three structures $\Omega_0, \Omega_1, \Omega_2$ in \mathcal{C} and embeddings $e : \Omega_0 \rightarrow \Omega_1$ and $f : \Omega_0 \rightarrow \Omega_2$ there exists $\Omega_3 \in \mathcal{C}$ and embeddings $e' : \Omega_1 \rightarrow \Omega_3$ and $f' : \Omega_2 \rightarrow \Omega_3$ such that $e' \circ e = f' \circ f$. For any Fraïssé class \mathcal{C} there exists an up to isomorphism unique homogeneous structure $\Delta_{\mathcal{C}}$, called the Fraïssé limit of \mathcal{C} , whose *age*, i.e., the class of its finitely generated substructures up to isomorphism, equals \mathcal{C} . Conversely, the age of any homogeneous structure in a countable language is a Fraïssé class.

For example, the *random graph* $G = (V, E)$ is the Fraïssé limit of the class of finite undirected graphs without loops, and similarly there exist a *random partial order*, a *random tournament*, *random hypergraphs*, a *random digraph*, and so forth. Let us stick to the first example for a moment and let us define a class of computational problems about finite graphs as follows. Call quantifier-free formulas in the language of graphs *graph formulas*. A graph formula $\Phi(x_1, \dots, x_m)$ is *satisfiable in a graph* iff there exists a graph H and an m -tuple a of elements in H such that $\Phi(a)$ holds in H . Now let $\Psi = \{\psi_1, \dots, \psi_n\}$ be a finite set of graph formulas. Then Ψ gives rise to the following computational problem.

Graph-SAT(Ψ)

INSTANCE: A set of variables W and a graph formula of the form $\Phi = \phi_1 \wedge \dots \wedge \phi_l$ where each ϕ_i for $1 \leq i \leq l$ is obtained from one of the formulas ψ in Ψ by substituting the variables from ψ by variables from W .

QUESTION: Is Φ satisfiable in a graph?

In words, an instance of Graph-SAT(Ψ) asks whether there exists a (finite) graph with satisfies a conjunction of properties; which properties can appear is restricted by the fixed set of graph formulas Ψ . Therefore, the computational complexity increases with Ψ in the sense that $\Psi \subseteq \Psi'$, then any algorithm for Graph-SAT(Ψ') solves Graph-SAT(Ψ). It is easy to see that each problem Graph-SAT(Ψ) is in NP, i.e., solvable in nondeterministic polynomial time.

The connection with CSPs is that every problem $\text{Graph-SAT}(\Psi)$ can be translated into $\text{CSP}(\Gamma_\Psi)$ for a finite language reduct Γ_Ψ of the random graph $G = (V, E)$ and vice-versa. For one direction, let $\Psi = \{\psi_1, \dots, \psi_n\}$ be a set of graph formulas. To this set we assign a reduct Γ_Ψ of G which has for each ψ_i a relation R_i consisting of those tuples of elements of V that satisfy ψ_i (where the arity of R_i is given by the number of distinct variables that occur in ψ_i). One readily sees that any algorithm for $\text{Graph-SAT}(\Psi)$ can be adapted to $\text{CSP}(\Gamma_\Psi)$ and vice-versa, and so the problems are essentially the same. For the other direction, if Γ is a reduct of G in a finite language, then each of its relations is defined by a first-order formula over G , and indeed even by a quantifier-free first-order formula (i.e., a graph formula), since homogeneity and ω -categoricity imply quantifier elimination. Let Ψ_Γ the set of those graph formulas. Again, one easily checks that $\text{Graph-SAT}(\Psi_\Gamma)$ and $\text{CSP}(\Gamma)$ are basically the same problem [13].

Now let \mathcal{C} be an arbitrary Fraïssé class of finitely generated structures in a finite language. As in the case of graphs, for a finite set Ψ of quantifier-free first-order formulas in the language of \mathcal{C} , we can define the following computational problem.

$\mathcal{C}\text{-SAT}(\Psi)$

INSTANCE: A set of variables W and a formula of the form $\Phi = \phi_1 \wedge \dots \wedge \phi_l$ where each ϕ_i for $1 \leq i \leq l$ is obtained from one of the formulas ψ in Ψ by substituting the variables from ψ by variables from W .

QUESTION: Is Φ satisfiable in a structure in \mathcal{C} ?

As before, each problem $\mathcal{C}\text{-SAT}(\Psi)$ is equivalent to $\text{CSP}(\Gamma_\Psi)$ for an appropriate reduct Γ_Ψ of the Fraïssé limit $\Delta_{\mathcal{C}}$ of \mathcal{C} and vice-versa. Hence, classifying the complexity of the problems $\mathcal{C}\text{-SAT}(\Psi)$ and classifying the complexity of the constraint satisfaction problems of reducts of $\Delta_{\mathcal{C}}$ is one and the same thing. Complete classifications have been obtained so far for the following countable homogeneous structures.

- the empty structure $(\mathbb{N}, =)$ [7];
- the order of the rationals (\mathbb{Q}, \leq) [9];
- the random graph $G = (V, E)$ [13].

In each of the three cases the classifications resulted in dichotomies: the CSPs of the reducts turned out to be either NP-complete or in P (i.e., solvable in polynomial time, which we will henceforth refer to as *tractable*). While there exist CSPs of homogeneous digraphs which are undecidable [18], the following representability condition for a Fraïssé class \mathcal{C} , arguably reasonable for the most interesting computational problems, forces $\mathcal{C}\text{-SAT}$ problems to be in NP, and could possibly imply a general dichotomy. Let τ be a finite relational signature. A class \mathcal{C} of finite τ -structures is called *finitely bounded* iff there exists a finite set of finite τ -structures \mathcal{F} such that \mathcal{C} consists precisely of those finite τ -structures which do not embed any element of \mathcal{F} . A relational structure is called finitely bounded iff its age is finitely bounded. When Γ is a finite language reduct of a finitely bounded homogeneous structure Δ , then $\text{CSP}(\Gamma)$ is easily seen to be in NP. We conjecture the following.

Conjecture 5. *Let Δ be a finitely bounded homogeneous structure, and let Γ be a finite language reduct of Δ . Then $\text{CSP}(\Gamma)$ is either in P or NP-complete.*

The advantage of translating $\mathcal{C}\text{-SAT}$ problems into constraint satisfaction problems of reducts of the Fraïssé limit $\Delta_{\mathcal{C}}$ is that it allows for the algebraic approach via clones, as we will now outline.

Firstly, by Theorem 3 we have that if Γ, Γ' are ω -categorical structures on the same domain and $\text{Pol}(\Gamma) \supseteq \text{Pol}(\Gamma')$, then $\text{CSP}(\Gamma)$ has a polynomial-time reduction to $\text{CSP}(\Gamma')$. On a theoretical level, this implies that when one wants to classify the complexity of the CSPs of all reducts of $\Delta_{\mathcal{C}}$, it suffices to consider all polymorphism clones of reducts of $\Delta_{\mathcal{C}}$ – reducts with equal polymorphism clones are polynomial-time equivalent. Those clones are precisely the closed function clones which contain the automorphism group $\text{Aut}(\Delta_{\mathcal{C}})$ of $\Delta_{\mathcal{C}}$. Moreover, if a closed function clone corresponds to a tractable (i.e., polynomial-time solvable) CSP, then so do all closed function clones containing it; if it corresponds to a NP-hard CSP, then so do all closed function clones above $\text{Aut}(\Delta_{\mathcal{C}})$ contained in it. In the case of an existing dichotomy, one thus has to find the border between tractability and NP-hardness in the lattice of closed function clones containing $\text{Aut}(\Delta_{\mathcal{C}})$. On a practical level, this implies that if the CSP of a reduct Γ is in P, then this is witnessed by the presence of certain functions in $\text{Pol}(\Gamma)$. And indeed, the presence of polymorphisms with certain properties have been successfully translated into algorithms in the classifications above – see [12].

The second use of polymorphism clones is that Theorem 4 allows us to compare CSPs on different domains, resulting in a tool both for showing tractability as well as for showing hardness. As for the latter, it turns out to be convenient to show NP-hardness of $\text{CSP}(\Gamma)$ by exposing a continuous clone homomorphism from $\text{Pol}(\Gamma)$ onto $\text{Pol}(\Gamma')$, where Γ' is an ω -categorical or finite structure with a hard CSP [11]. In practice, Γ' will generally be finite; in fact, Γ' will often be any structure on a two-element set with a *trivial* polymorphism clone, i.e., the clone $\mathbf{1}$ of all projections on a two-element set (which is the polymorphism clone of NP-complete structures). Since $\mathbf{1}$ is isomorphic to the smallest function clone (i.e., the clone of projections) on any finite set with at least two elements, and on finite domain smaller polymorphism clones correspond to harder CSPs, a continuous clone homomorphism to $\mathbf{1}$ is in a sense the strongest finite reason for NP-hardness (more precisely, it implies that Γ *pp-interprets* all finite structures [11]; see also [17]).

Notation 6. We write $\mathbf{1}$ for the clone of all projections on a two-element set.

It is an open conjecture, and indeed the main conjecture for CSPs of finite structures known as the *tractability conjecture*, that under a cosmetic assumption (the assumption of having an *idempotent* polymorphism clone, see Section 3) on a finite structure Γ , the CSP for Γ is NP-complete if there exists a clone homomorphism from $\text{Pol}(\Gamma)$ onto $\mathbf{1}$, and in P otherwise. Clearly, if there exists such a homomorphism, then $\text{CSP}(\Gamma)$ is NP-complete; the open part is the other direction. Note that if there is no such homomorphism, then this is witnessed by equations which hold in $\text{Pol}(\Gamma)$ but which cannot be satisfied in $\mathbf{1}$. Numerous equations have been translated into algorithms, and indeed every non-trivial set of equations of an idempotent clone translates into an algorithm if one believes in the tractability conjecture. In the ω -categorical setting, we cannot purely rely on equations, but need to take into account the topology on the functions – at least with what we know today. Recently, research has been conducted investigating the role of this topology [17], and about how to show tractability in case there exists no continuous homomorphism to $\mathbf{1}$.

A Ramsey-theoretic method. The behavior of polymorphism clones of reducts of homogeneous structures Δ in a finite relational language seems to be particularly close to that of function clones on finite sets when Δ satisfies a particular combinatorial property. Let τ be a relational signature. We say that a class \mathcal{C} of finite τ -structures is a *Ramsey class* (in the sense of [37]) iff for all $\Omega_0, \Omega_1 \in \mathcal{C}$ there exists $\Omega_2 \in \mathcal{C}$ such that for all colorings of the

copies of Ω_0 in Ω_2 with two colors there exists an isomorphic copy Ω'_1 of Ω_1 in Ω_2 such that all copies of Ω_0 in Ω'_1 have the same color. A relational structure is called *Ramsey* iff its age is a Ramsey class. When \mathcal{C} is a relational Fraïssé class and $\Delta_{\mathcal{C}}$ its Fraïssé limit, then it is equivalent to call $\Delta_{\mathcal{C}}$ Ramsey iff for all $\Omega_0, \Omega_1 \in \mathcal{C}$ and all colorings of the copies of Ω_0 in $\Delta_{\mathcal{C}}$ there exists a copy of Ω_1 in $\Delta_{\mathcal{C}}$ on which the coloring is constant. Examples of Fraïssé classes which are Ramsey classes are the class of finite ordered undirected graphs, the class of finite linear orders, and the class of finite partial orders with a linear extension [39, 38, 1].

We remark that, for example, neither the random graph nor the random partial order are Ramsey, and thus seem to fall out of this framework. However, they are themselves reducts of homogeneous Ramsey structures, namely the random ordered graph (i.e., the Fraïssé limit of the class of finite ordered undirected graphs) and the random partial order with a random linear extension (i.e., the Fraïssé limit of the class of finite partial orders with a linear extension).

The Ramsey property can be exploited as follows. Let Ξ be a structure. The *type* of a tuple $b = (b^1, \dots, b^n)$ of elements of Ξ , denoted by $\text{tp}(b)$, is the set of first-order formulas $\phi(x_1, \dots, x_n)$ such that $\phi(b^1, \dots, b^n)$ holds in Ξ . Now let Ξ_1, \dots, Ξ_m be structures. For an element a of the product $\Xi_1 \times \dots \times \Xi_m$ and $1 \leq i \leq m$, we write a_i for the i -th coordinate of a . The *type* of a tuple (a^1, \dots, a^n) of elements $a^1, \dots, a^n \in \Xi_1 \times \dots \times \Xi_m$, denoted by $\text{tp}(a^1, \dots, a^n)$, is the m -tuple containing the types of (a_i^1, \dots, a_i^n) in Ξ_i for each $1 \leq i \leq m$. A function $f : \Xi_1 \times \dots \times \Xi_m \rightarrow \Omega$ is called *canonical* iff it sends finite tuples of equal type in $\Xi_1 \times \dots \times \Xi_m$ to tuples of equal type in Ω ; that is, whenever $\text{tp}(a^1, \dots, a^n) = \text{tp}(b^1, \dots, b^n)$, then $\text{tp}(f(a^1), \dots, f(a^n)) = \text{tp}(f(b^1), \dots, f(b^n))$. For a relational structure Δ and elements c_1, \dots, c_n of Δ , we write $(\Delta, c_1, \dots, c_n)$ for the expansion of Δ by the constants c_1, \dots, c_n . The structure Δ is *ordered* iff it has a linear order among its relations. Now the following holds [18, 12].

Theorem 7 (Bodirsky, Pinsker and Tsankov [18]). *Let Δ be an ordered homogeneous Ramsey structure in a finite relational language, and let $\mathcal{C} \supseteq \text{Aut}(\Delta)$ be a closed function clone. Then for all $f \in \mathcal{C}$ and all $c_1, \dots, c_n \in \Delta$ there exists a function $g \in \mathcal{C}$ which is canonical as a function on $(\Delta, c_1, \dots, c_n)$, and which agrees with f on $\{c_1, \dots, c_n\}$.*

Thus under these conditions on Δ , if there is a function f in a polymorphism clone of a reduct of Δ which does something of interest (e.g., algorithmically) on a finite set $\{c_1, \dots, c_n\}$, then there is also a canonical function in this clone which does the same. Note that canonical functions on $(\Delta, c_1, \dots, c_n)$ are finite objects in the following sense. Every canonical function $f : (\Delta, c_1, \dots, c_n)^m \rightarrow (\Delta, c_1, \dots, c_n)$ defines an m -ary function $T(f)$ on the types of $(\Delta, c_1, \dots, c_n)$ in an obvious way, by the very definition of canonicity. Moreover, this type function $T(f)$ determines f in the sense that if two canonical functions $f, g : (\Delta, c_1, \dots, c_n)^m \rightarrow (\Delta, c_1, \dots, c_n)$ have identical type functions $T(f) = T(g)$, then any closed function clone containing $\text{Aut}(\Delta)$ contains f iff it contains g . Since $(\Delta, c_1, \dots, c_n)$ is homogeneous in a finite language, the type functions are finite objects: $T(f)$ is completely determined by its values on the types of tuples of length q , where q is the maximal arity of a relation in $(\Delta, c_1, \dots, c_n)$; moreover, there are only finitely many types of q -tuples since $(\Delta, c_1, \dots, c_n)$ is ω -categorical. As finite objects, these type functions can effectively be used in algorithms [13].

An example of an application of Theorem 7 is the following.

Theorem 8 (Bodirsky, Pinsker and Tsankov [18]). *Let Δ be an ordered homogeneous Ramsey structure in a finite relational language, and let Γ be a reduct in a finite language. Then there*

exist c_1, \dots, c_n , $m \geq 1$, and m -ary canonical functions f_1, \dots, f_k on $(\Delta, c_1, \dots, c_n)$ such that for all reducts Γ' we have that $\text{Pol}(\Gamma') \setminus \text{Pol}(\Gamma)$ is either empty or contains one of the functions f_1, \dots, f_k .

In words, under the above conditions the finite language reducts of Δ can be distinguished by functions which are canonical after adding finitely many constants to the language of Δ . If we assume that Δ is finitely bounded, which makes Δ in a way finitely representable, then Theorem 8 can even be implemented in an algorithm. This yields the following effective variant of the theorem.

Theorem 9 (Bodirsky, Pinsker and Tsankov [18]). *Let Δ be an ordered homogeneous Ramsey structure which is finitely bounded, and let Γ, Γ' be finite language reducts of Δ . Then the problem whether or not $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma')$, where the relations of Γ and Γ' are given by quantifier-free formulas over Δ , is decidable.*

This gives hope that tractability of CSPs of reducts of Δ is captured by the canonical functions in their polymorphism clones – cf. Section 3.

The modern proof of Theorem 8 (yet unpublished but available on request) is based on a beautiful characterization of the Ramsey property for homogeneous structures which links Ramsey theory with topological dynamics [32]. A topological group \mathcal{G} is called *extremely amenable* iff whenever it acts continuously on a compact Hausdorff topological space X , then this action has a fixed point, i.e., there exists $x \in X$ such that $g(x) = x$ for all $g \in \mathcal{G}$. Let Δ be an ordered homogeneous relational structure. Then Δ is Ramsey iff $\text{Aut}(\Delta)$, viewed as an abstract topological group, is extremely amenable.

3. OPEN PROBLEMS

The research questions presented here are all related to Conjecture 5 in one way or another; in fact, one can put them together so that they constitute a systematic program for proving the conjecture. After stating the questions, in Section 4, I will discuss the questions in the context of the conjecture. I emphasize, however, that each of them has its own mathematical value independently of the truth of the conjecture.

The first set of questions concerns the connection between the algebraic and the topological structure of clones in the light of Theorem 2. As for the link to constraint satisfaction, recall that for finite structures Γ , the complexity of $\text{CSP}(\Gamma)$ only depends on the algebraic structure of $\text{Pol}(\Gamma)$, whereas in the ω -categorical setting, one also has to take the topology on $\text{Pol}(\Gamma)$ into consideration (Theorem 4). A first question, which we shall then refine, is the following.

Question 10. *Are there conditions on oligomorphic algebras under which we can drop the continuity condition in Theorem 2?*

One such condition is to allow only very simple algebras \mathfrak{B} . Of particular interest are, of course, continuous clone homomorphisms to $\mathbf{1}$ (i.e., term clones of algebras all of whose functions are projections), since they are the major source of hardness proofs. In fact, we believe that for certain structures they are the *unique* source of hardness proofs: as already mentioned, the finite tractability conjecture states that if a finite structure Γ satisfies a cosmetic condition, then $\text{CSP}(\Gamma)$ is NP-hard if and only if there exists a homomorphism from $\text{Pol}(\Gamma)$ to $\mathbf{1}$. That condition, which requires that $\text{Pol}(\Gamma)$ is *idempotent*, i.e., all $f \in \text{Pol}(\Gamma)$ satisfy the equation $f(x, \dots, x) = x$, can always be assumed: every CSP of a finite structure is equivalent to a CSP of a finite structure with an idempotent polymorphism clone [21]. This

fact can be derived in two steps: first, one shows that Γ can be assumed to be a *core*, i.e., all endomorphisms of Γ are automorphisms. One then shows that it is possible to add finitely many constants to the language of Γ without increasing the complexity of its CSP – adding one constant for each element of the domain, this forces all polymorphisms to be idempotent. In the ω -categorical setting, one cannot simply assume that $\text{Pol}(\Gamma)$ be idempotent; indeed, it would then certainly fail to be oligomorphic, containing no unary functions at all except the identity. However, it is possible to perform an analog of the first step, and assume that Γ is a *model-complete core*, meaning that $\text{Aut}(\Gamma)$ is (topologically) dense in the endomorphism monoid of Γ (i.e., every endomorphism *locally* looks like an automorphism) [4]. Moreover, as in the finite case, adding finitely many constants to the language of Γ then does not increase the complexity of its CSP [4]. Until now, similarly to the situation for finite templates, we do not know of a reduct of a finitely bounded homogeneous structure which is a model-complete core with an NP-hard CSP, but which has no continuous homomorphism to $\mathbf{1}$ after adding finitely many constants to the language.

Back to the continuity condition, we therefore ask the following.

Question 11. *If a closed oligomorphic clone has a clone homomorphism to $\mathbf{1}$ (i.e., it satisfies no non-trivial equations), does it always have a continuous clone homomorphism to $\mathbf{1}$? If not, are there further conditions on the clone (model complete core etc.) which imply a positive answer?*

There exists considerable literature about automorphism groups of ω -categorical structures Γ which are *reconstructible*, i.e., where the topology on $\text{Aut}(\Gamma)$ is uniquely determined by the algebraic group structure; this is for instance the case when $\text{Aut}(\Gamma)$ has the so-called *small index property*, that is, all subgroups of countable index are open. The small index property has for instance been shown for $\text{Aut}(\mathbb{N}, =)$ [24]; for $\text{Aut}(\mathbb{Q}; <)$ and for the automorphism group of the atomless Boolean algebra [42]; the automorphism group of the random graph [29]; for all ω -categorical ω -stable structures [29]; for the automorphism groups of the Henson graphs [27]. The notion of reconstruction makes perfect sense for function clones, and is of importance for our purposes. Call a closed function clone *reconstructible* iff all isomorphisms with other closed function clones are homeomorphisms. Recent research has shown that for some homogeneous ω -categorical structures with a reconstructible automorphism group, the reconstructibility carries over to the polymorphism clone of the structure [16].

Question 12. *Let \mathcal{C} be an oligomorphic polymorphism clone whose group of invertible unary functions is reconstructible. When can we conclude that \mathcal{C} is reconstructible as well?*

We remark that there exists an example of two ω -categorical structures whose automorphism groups are isomorphic as groups but not as topological groups [25], and that this example has recently been expanded to polymorphism clones by David Evans in a yet unpublished note.

It is well-known that every Baire measurable homomorphism between Polish groups is continuous (see e.g. [31]). So let us remark that there exists a model of ZF+DC where every set is Baire measurable [41]. For the structures Γ that we need to model computational problems as $\text{CSP}(\Gamma)$ it therefore seems fair to assume that the abstract algebraic structure of $\text{Aut}(\Gamma)$ always determines its topological structure; consistency of this statement with ZF has already been observed in [34]. Hence, one could hope to find a model of ZF in which polymorphism clones of ω -categorical structures are reconstructible, or in which all homomorphisms of such clones to $\mathbf{1}$ are continuous.

Question 13. *Do oligomorphic polymorphism clones have reconstruction in an appropriate model of ZF? Are all homomorphisms from oligomorphic polymorphism clones to $\mathbf{1}$ continuous in an appropriate model of ZF?*

The concept of a *topological clone* appeared as a necessity for formulating Theorem 2; it is indeed natural considering the importance of abstract clones (known in disguise as *varieties*) for universal algebra and the natural presence of topological groups in various fields of mathematics, in particular topological dynamics. It is known that the closed permutation groups on a countable set are precisely those topological groups that are Polish and have a left-invariant ultrametric [2].

Question 14. *Which topological clones appear as closed function clones on a countably infinite set?*

We now turn to the study of polymorphism clones of reducts of homogeneous Ramsey structures. Here, the approach via canonical functions, based on Theorem 7 and the idea that we keep sufficient information about a clone when we add a sufficiently large finite number of constants and then only consider its canonical functions, has proven extremely fruitful. For example, this was the strategy in the Graph-SAT dichotomy classification [13], and in many other applications [40, 15, 18, 12, 14]; confer also Theorems 8 and 9. Generalizing the Graph-SAT strategy, we arrive at the following ideas.

In the following, let Γ be a reduct of a finitely bounded ordered homogeneous Ramsey structure Δ , and let $c_1, \dots, c_n \in \Delta$. Then the set of finitary canonical functions on $(\Delta, c_1, \dots, c_n)$ forms a closed function clone, and hence so does the intersection of this clone with $\text{Pol}(\Gamma)$, which we call the *canonical fragment* of $\text{Pol}(\Gamma)$ with respect to c_1, \dots, c_n . By Theorem 8, this canonical fragment still contains considerable information about $\text{Pol}(\Gamma)$. Recall that its functions define functions on the types of $(\Delta, c_1, \dots, c_n)$; in fact, these “type functions” form a clone on a finite set. We call this clone the *type clone* of $\text{Pol}(\Gamma)$ with respect to c_1, \dots, c_n , and denote it by $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$. There are infinitely many choices for c_1, \dots, c_n , but up to type equivalence, only finitely many for each n since Γ is ω -categorical [28]. In the Graph-SAT dichotomy, these type clones happened to capture the computational complexity of $\text{CSP}(\Gamma)$ [13].

Question 15. *Does the complexity of $\text{CSP}(\Gamma)$ only depend on the algebraic structure of its type clones?*

More precisely, in the Graph-SAT classification it turned out that when $\text{CSP}(\Gamma)$ is tractable, then this fact was captured by some canonical polymorphism, which provided the algorithm; in particular, the answer to the following question was positive. It is nourished by the belief that if $\text{Pol}(\Gamma)$ contains a function which implies tractability, and therefore is of use in some algorithm, then this function can be “canonized” and therefore appears in some type clone (cf. Theorem 7).

Question 16. *If Γ is tractable, are there necessarily $c_1, \dots, c_n \in \Delta$ such that $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$ corresponds to a tractable CSP?*

The following question asks about the converse. The intuition behind it, again true in the Graph-SAT case, is that algorithms for the type clone can be “lifted” back to the original clone.

Question 17. *If there exist $c_1, \dots, c_n \in \Delta$ such that $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$ is tractable, is $\text{Pol}(\Gamma)$ tractable?*

The following can be seen as a complexity-free variant of the preceding two questions. It is known that if $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$ satisfies non-trivial equations for some $c_1, \dots, c_n \in \Delta$, then so does $\text{Pol}(\Gamma)$ [17]. Therefore, if some type clone corresponds to a tractable, and if the finite tractability conjecture holds, then $\text{Pol}(\Gamma)$ does not have a homomorphism to $\mathbf{1}$ (and hence, if NP-hard, would have to have another source of hardness). We do not know if the converse is true as well.

Question 18. *If $\text{Pol}(\Gamma)$ does not have a homomorphism to $\mathbf{1}$, are there $c_1, \dots, c_n \in \Delta$ such that $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$ has no homomorphism to $\mathbf{1}$?*

We often think of the Ramsey property as an additional property of Fraïssé classes (and indeed, most homogeneous structures are not Ramsey); we do require the property for our methods. However, we do not require the reducts Γ , but only some “base structure” Δ in which they are definable, to be Ramsey. Hence, if the answer to the following question were positive, then our methods would work for all homogeneous structures.

Question 19. *Can every finitely bounded Fraïssé class be extended by finitely many relations to a finitely bounded Fraïssé class which is in addition Ramsey?*

We remark that this question has recently received considerable attention, and in particular the answer is positive for all Fraïssé classes of digraphs – see [33] for further references.

Let us turn to concrete classes of CSPs. Studying those has its own interest, just like the dichotomies of Graph-SAT problems [13] and temporal constraints [9, 8]; moreover they provide sources of examples for the general questions.

The class of finite partial orders is one of the most natural Fraïssé classes, and the answer to the following question would subsume some older results in theoretical computer science, e.g. in [19], and the classification [9]. The basis for a successful complexity classification of Poset-SAT problems has been established very recently in the form of the classification of the closed supergroups of the automorphism group of the random partial order [40].

Question 20. *Classify the complexity of $\text{CSP}(\Gamma)$, for all finite language reducts Γ of the random partial order. In other words, classify the complexity of Poset-SAT problems.*

The solution to the following would subsume a considerable amount of results in the literature, e.g., completely the papers [30], [6], and some results in [19], [35]. It would moreover require the extension of our methods to *functional* signatures, a venture interesting in itself.

Question 21. *Classify the complexity of $\text{CSP}(\Gamma)$, for all finite language reducts Γ of the atomless (= random) Boolean algebra.*

4. THE INFINITE TRACTABILITY CONJECTURE

Each of the above questions has its own interest for the understanding of oligomorphic function clones, oligomorphic algebras, and their connections with constraint satisfaction. However, these questions really are part of a bigger program around Conjecture 5. To make this connection with Conjecture 5 more evident, let me show an example of how a proof of the conjecture could look like. Let Γ be a finite language reduct of a finitely bounded homogeneous structure Δ .

- Assume that $\text{Pol}(\Gamma)$ does not have a continuous clone homomorphism to $\mathbf{1}$ (otherwise, $\text{CSP}(\Gamma)$ is NP-hard by Theorem 4).

- If the answer to Question 11 is positive (possibly in an appropriate model of ZF, cf. Question 13), then $\text{Pol}(\Gamma)$ satisfies non-trivial equations.
- Assuming that Question 19 has a positive answer, we can assume that Δ is Ramsey.
- If Question 18 has a positive answer, then some type clone $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$ satisfies non-trivial equations.
- Assuming the tractability conjecture for finite templates, $T_{c_1, \dots, c_n}(\text{Pol}(\Gamma))$ corresponds to a tractable CSP.
- Assuming a positive answer to Question 17, $\text{CSP}(\Gamma)$ is then tractable.

REFERENCES

- [1] Fred G. Abramson and Leo Harrington. Models without indiscernibles. *Journal of Symbolic Logic*, 43(3):572–600, 1978.
- [2] Howard Becker and Alexander Kechris. *The Descriptive Set Theory of Polish Group Actions*. Number 232 in LMS Lecture Note Series. Cambridge University Press, 1996.
- [3] Garrett Birkhoff. On the structure of abstract algebras. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):433–454, 1935.
- [4] Manuel Bodirsky. Complexity classification in infinite-domain constraint satisfaction. Mémoire d’habilitation à diriger des recherches, Université Diderot – Paris 7. Available at arXiv:1201.0856, 2012.
- [5] Manuel Bodirsky and Hubert Chen. Oligomorphic clones. *Algebra Universalis*, 57(1):109–125, 2007.
- [6] Manuel Bodirsky and Martin Hils. Tractable set constraints. *Journal of Artificial Intelligence Research*, 45:731–759, 2012.
- [7] Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of Computer Science Russia (CSR’06).
- [8] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 29–38, 2008.
- [9] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2):1–41, 2009. An extended abstract appeared in the Proceedings of the Symposium on Theory of Computing (STOC’08).
- [10] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.
- [11] Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. *Transactions of the AMS*. To appear. Preprint arxiv.org/abs/1203.1876.
- [12] Manuel Bodirsky and Michael Pinsker. Reducts of Ramsey structures. *AMS Contemporary Mathematics*, vol. 558 (*Model Theoretic Methods in Finite Combinatorics*), pages 489–519, 2011.
- [13] Manuel Bodirsky and Michael Pinsker. Schaefer’s theorem for graphs. In *Proceedings of the Symposium on Theory of Computing (STOC)*, pages 655–664, 2011. Preprint of the full journal version available from arxiv.org/abs/1011.2894.
- [14] Manuel Bodirsky and Michael Pinsker. Minimal functions on the random graph. *Israel Journal of Mathematics*, 200(1):251–296, 2014.
- [15] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. The 42 reducts of the random ordered graph. Preprint arXiv:1309.2165, 2013.
- [16] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Reconstructing the topology of clones. Preprint arXiv:1312.7699, 2013.
- [17] Manuel Bodirsky, Michael Pinsker, and András Pongrácz. Projective clone homomorphisms. Preprint arXiv:1409.4601, 2014.
- [18] Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. *Journal of Symbolic Logic*. to appear. A conference version appeared in the Proceedings of LICS 2011, pages 321–328.
- [19] Mathias Broxvall and Peter Jonsson. Point algebras for temporal reasoning: Algorithms and complexity. *Artificial Intelligence*, 149(2):179–220, 2003.
- [20] Andrei A. Bulatov, Andrei A. Krokhin, and Peter Jeavons. Constraint satisfaction problems and finite algebras. In *Proceedings of ICALP*, pages 272–282, 2000.

- [21] Andrei A. Bulatov, Andrei A. Krokhin, and Peter G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.
- [22] Stanley N. Burris and Hanamantagouda P. Sankappanavar. *A Course in Universal Algebra*. Springer Verlag, Berlin, 1981.
- [23] Peter J. Cameron. *Oligomorphic permutation groups*. Cambridge University Press, Cambridge, 1990.
- [24] John Dixon, Peter M. Neumann, and Simon Thomas. Subgroups of small index in infinite symmetric groups. *Bulletin of the London Mathematical Society*, 18(6):580–586, 1986.
- [25] David M. Evans and Paul R. Hewitt. Counterexamples to a conjecture on relative categoricity. *Annals of Pure and Applied Logic*, 46(2):201–209, 1990.
- [26] Martin Goldstern and Michael Pinsker. A survey of clones on infinite sets. *Algebra Universalis*, 59:365–403, 2008.
- [27] Bernhard Herwig. Extending partial isomorphisms for the small index property of many ω -categorical structures. *Israel Journal of Mathematics*, 107:93–123, 1998.
- [28] Wilfrid Hodges. *Model theory*. Cambridge University Press, 1993.
- [29] Wilfried Hodges, Ian Hodkinson, Daniel Lascar, and Saharon Shelah. The small index property for ω -categorical ω -stable structures and for the random graph. *Journal of the London Mathematical Society*, 48(2):204–218, 1993.
- [30] Peter Jonsson and Thomas Drakengren. A complete classification of tractability in RCC-5. *Journal of Artificial Intelligence Research*, 6:211–221, 1997.
- [31] Alexander Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Texts in Mathematics*. Springer, 1995.
- [32] Alexander Kechris, Vladimir Pestov, and Stevo Todorcevic. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis*, 15(1):106–189, 2005.
- [33] Claude Laflamme, Jakub Jasiński, Lionel Nguyen Van Thé, and Robert Woodrow. Ramsey precompact expansions of homogeneous directed graphs. *Electronic Journal of Combinatorics*, 21(4), 2014. 31 pages.
- [34] Daniel Lascar. Autour de la propriété du petit indice. *Proceedings of the London Mathematical Society*, 62(1):25–53, 1991.
- [35] Kim Marriott and Martin Odersky. Negative Boolean constraints. *Theoretical Computer Science*, 160(1&2):365–380, 1996.
- [36] Dragan Mašulović and Maja Pech. Oligomorphic transformation monoids and homomorphism-homogeneous structures. *Fundamenta Mathematicae*, 212(1):17–34, 2011.
- [37] Jaroslav Nešetřil. Ramsey theory. *Handbook of Combinatorics*, pages 1331–1403, 1995.
- [38] Jaroslav Nešetřil and Vojtěch Rödl. Ramsey classes of set systems. *Journal of Combinatorial Theory, Series A*, 34(2):183–201, 1983.
- [39] Jaroslav Nešetřil and Vojtěch Rödl. The partite construction and Ramsey set systems. *Discrete Mathematics*, 75(1-3):327–334, 1989.
- [40] Péter Pál Pach, Michael Pinsker, Gabriella Pluhár, András Pongrácz, and Csaba Szabó. Reducts of the random partial order. *Advances in Mathematics*, 267:94–120, 2014.
- [41] Saharon Shelah. Can you take Solovay’s inaccessible away? *Israel Journal of Mathematics*, 48(1):1–47, 1984.
- [42] John K. Truss. Infinite permutation groups. II. Subgroups of small index. *Journal of Algebra*, 120(2):494–515, 1989.